
Adaptive Mesh Coarsening for Quadrilateral and Hexahedral Meshes

Jason F. Shepherd¹, Mark W. Dewey², Adam C. Woodbury³, Steven E. Benzley³, Matthew L. Staten^{1,4}, and Steven J. Owen¹

¹ Sandia National Laboratories

jfsheph@sandia.gov, mlstate@sandia.gov, sjowen@sandia.gov

² Computational Simulation Software, Inc. markwdewey@gmail.com

³ Brigham Young University aw18@byu.net, seb@et.byu.edu

⁴ Carnegie Mellon University mstaten@cmu.edu

Abstract: Mesh adaptation methods can improve the efficiency and accuracy of solutions to computational modeling problems. In many applications involving quadrilateral and hexahedral meshes, local modifications which maintain the original element type are desired. For triangle and tetrahedral meshes, effective refinement and coarsening methods that satisfy these criteria are available. Refinement methods for quadrilateral and hexahedral meshes are also available. However, due to the added complexity of maintaining and satisfying constraints in quadrilateral and hexahedral mesh topology, little research has occurred in the area of coarsening or simplification. This paper presents methods to locally coarsen conforming all-quadrilateral and all-hexahedral meshes. The methods presented provide coarsening while maintaining conforming all-quadrilateral and all-hexahedral meshes. Additionally, the coarsening is not dependent on reversing a previous refinement. Several examples showing localized coarsening are provided.

Keywords: Quadrilateral, Hexahedral, Mesh, Coarsening, Simplification, Adaptivity, Refinement

1 Introduction

Mesh adaptation methods can improve the efficiency and accuracy of solutions to computational modeling problems. For a given model, there are usually regions that require greater mesh density than others to improve solution efficiency, reduce error or uncertainty in high gradient regions, or more accurately represent the model geometry. Regions where high accuracy is not critical or where gradients are low can generally be modeled with lower mesh density. Since the computational time required in a finite element analysis

is directly related to the number of elements in the model being analyzed, it is advantageous to produce a mesh that has as few elements as possible. Therefore, in an ideal analysis, each region in the model should have enough elements to produce a good solution, but no more.

Due to the complexity inherent in many mesh generation algorithms, it is often difficult to create an initial mesh that optimizes both accuracy and efficiency. Although some control over mesh density is possible, an initial mesh will almost always contain regions that have too few elements, regions that have too many elements, or both. In addition, some applications require mesh density to evolve throughout an analysis as areas of high and low activity change with time [1, 2, 3, 4]. For these reasons, much research has been devoted to the development of mesh modification tools that make it possible to adjust element density in specific regions either before or during analysis.

Mesh adaptation consists of both refinement and coarsening. Refinement is the process of adding elements to a mesh while coarsening is the process of removing elements from a mesh. By refining areas that have too few elements and coarsening areas that have too many elements, a more accurate and efficient analysis can be performed. Mesh adaptation methods are also useful in visual applications where objects far from view can be highly simplified while objects closer to view should have more detail. Because computer visualizations are typically embedded on a mesh, efficient algorithms for mesh adaptation are valuable for improving memory performance for views consisting of large numbers of mesh elements.

To date, most of the research in mesh adaptation has focused on refinement techniques for increasing local element density [5, 6]. Complementary algorithms for decreasing local element density by element removal (i.e., coarsening) could be a powerful companion tool to refinement algorithms, potentially allowing more flexible mesh adaptation. For example, given a uniform mesh, the mesh density in an area of interest may be increased by established refinement techniques and decreased away from the areas of interest using a coarsening technique. Rather than remeshing the model, the base mesh may be modified using refinement and coarsening tools; this would permit increased resolution and accuracy in the results while maintaining a similar computation time for the entire model. Furthermore, a given model may require adaptation in different locations depending on different load cases, adaptation by both refinement and coarsening from a single base mesh may allow more efficient and robust generation of meshes appropriate for varied circumstances. In spite of its potential benefits, coarsening is an area of research which has received limited attention.

In this paper we describe algorithms for performing coarsening on all-quadrilateral and all-hexahedral meshes while maintaining conforming mesh topology through the coarsening process to prevent the creation of non-quadrilateral or non-hexahedral elements. In the following sections we will discuss related work in mesh coarsening or simplification, outline our algorithms and demonstrate the algorithms on several examples.

2 Background

Mesh adaptation is a field which has received extensive study among both computational mechanics and computer graphics researchers. Generally these two fields have not collaborated due to the many additional restrictions associated with computational mechanics but unnecessary in computer graphics. One example of these additional restrictions in computational mechanics is that a mesh must accurately represent the model geometry by ensuring that the nodes representing a curve or surface of the model do not move off the geometry, whereas in graphics a sufficiently low level of detail might justify combining surfaces and/or curves.

To effectively achieve the objectives of mesh adaptation, a truly general quadrilateral/hexahedral coarsening algorithm should:

1. Preserve a conforming all-hexahedral or all-quadrilateral mesh
2. Restrict mesh topology and density changes to defined regions
3. Work on both structured and unstructured meshes
4. Not be limited to only undoing previous refinement

2.1 Triangle and Quadrilateral Simplification Algorithms

Triangular meshes in computer graphics and computational mechanics are common due to the relative simplicity of generating the meshes from these simplex elements. Triangle meshing algorithms are well-established and ongoing efforts in the research community continue to improve the quality of these meshes. Triangle mesh simplification algorithms begin with an existing base mesh, consisting of triangles, and modify the topology to remove triangles, improve quality and/or geometric integrity. A survey of triangle mesh coarsening algorithms is documented by Cignoni, et al. [7], highlighting the major simplification methodologies, including coplanar facet merging, controlled vertex/edge/face decimation, retiling, energy function optimization, vertex clustering, wavelet based approaches, and simplification via intermediate hierarchical representation. Additional surveys that compare smaller sets of algorithms are also given in [8].

One of the foremost algorithms of triangle mesh simplification was developed by Garland, et al. [9]. The approach is fast, reliable, and is also generally applicable to any polygon mesh. The algorithm assumes that the mesh is composed entirely of triangles, or can be broken into a mesh composed of triangles. It is designed to combine surfaces and curves that are indistinguishable when rendered at a low level of detail. Hoppe, et al. [10], demonstrate mesh adaptation respecting geometric curves and surfaces in order to preserve sharp corners and edges in the mesh representation.

While triangle meshes have widespread use, quadrilateral meshes are sometimes preferred in computational analysis due to some beneficial mathematical properties of the quadrilateral element that can result in increased solution

accuracy with fewer elements than triangle meshes [11]. Unfortunately, despite the wide availability of triangle mesh adaptation algorithms, most of algorithms developed for triangle mesh simplification cannot be adapted for use on quadrilateral meshes.

A number of efforts have been utilized for quadrilateral coarsening of structured meshes. Takeuchi, et al. [12], modified the approach developed by Garland, et al. [9], to simplify quadrilateral meshes; however, the process is designed for full-model simplification and may produce degenerate elements (i.e. quadrilaterals which are inverted or concave). Cheng, et al., [13] developed a method of coarsening a structured, all-quadrilateral mesh specifically for use on auto-body parts; however, this method has not been adapted for use in unstructured meshes. Kwak, et al., [14] performs simplification using remeshing algorithms; however, this global approach can be slow when only local adaptation is needed. Choi [15] describes an algorithm which can be used to undo previous refinement on both quadrilateral and hexahedral meshes; however, the reliance on knowledge of previous refinement restricts the algorithm from being used on a base mesh that has not been refined. Nikishkov [16] developed a quadtree method for mesh adaptation that allows both refinement and coarsening; however, his method requires the use of special elements or produces non-conforming elements.

2.2 Hexahedral Coarsening

Although hexahedral coarsening has been utilized in some modeling applications, no single algorithm has been developed that satisfies all the criteria listed above. This is, in large part, due to the topology constraints that exist in a conforming all-hexahedral mesh. These constraints make it difficult to modify mesh density without causing topology changes to propagate beyond the boundaries of a defined region [17, 18].

Since current hexahedral coarsening methods are unable to satisfy all the requirements listed above, they have limited application. For example, to prevent global topology changes, some algorithms introduce non-conforming or non-hexahedral elements into the mesh [1, 2, 19, 20, 21]. While this is a valid solution for some types of analysis, not all finite element solvers can accommodate hanging nodes or hybrid meshes. Other algorithms maintain a conforming all-hexahedral mesh, but they generally require either global topology changes beyond the defined coarsening region [18, 22, 23], structured mesh topology where predetermined transition templates can be used [24, 22], or prior refinement that can be undone [2, 19, 20]. These weaknesses severely limit the effectiveness of these algorithms on most real-world models.

3 Dual Methods

In recent years, a greater understanding of quadrilateral and hexahedral mesh topology has led to the development of many new quadrilateral and hexahe-

dral mesh operations [25, 26, 27, 28]. The algorithms presented in this paper utilize the dual representation of a quadrilateral/hexahedral mesh. In this section, we discuss dual-based operations which are useful for modification of quadrilateral and hexahedral meshes.

3.1 Quadrilateral Dual Methods

A dual chord is a set of quadrilaterals connected through pairs of opposite edges that extend through a mesh to connect back at the original starting edge (on a closed surface) or terminate at the mesh boundary (for a bounded quadrilateral mesh). In Figure 1, the dashed line highlights a single chord of the quadrilateral mesh shown.




Fig. 1. One method for removing a quadrilateral from a mesh to produce a conforming quadrilateral mesh requires simultaneously removing one of the chords associated with the quadrilateral.

Borden, et al. [23], illustrated that it is possible to remove an entire chord from a quadrilateral mesh, maintaining conformal connectivity, by simply collapsing the defining edges of the chord as shown in Figure 1. The removal of a chord reduces the number of quadrilaterals in the mesh and coarsens the quadrilaterals adjacent to the chord. Unfortunately, the effect of the coarsening extends along the entire length of the chord, which typically extends beyond the boundaries of a localized coarsening region. Benzley, et al. [24], extended the research of Borden, et al., by attempting to localize the chords to the coarsening region. However, their localization is dependent upon finding special element configurations in the original mesh.

Staten, et al. [29], extended this work eliminating the need for special element configurations. Simple chord operations (i.e. alterations to the mesh which change the connectivity of chords) are performed at the intersections of the chords in the coarsening region to alter the topology and create a single chord localized to the coarsening region. This chord can then be extracted coarsening only the localized region. The operations used to create the localized chord included the edge swap, face close and doublet insertion operations. Figure 2 shows an example of the chord operations utilized to modify a quadrilateral mesh to produce localized chords to the coarsening region.




Fig. 2. Chord operations: top right, Face Close; bottom left, Doublet Insertion; bottom right, Edge Swap.

Figure 3 demonstrates an example of the process outlined by Staten, et al., to locally coarsen a region of the mesh. In the second image from the left of Figure 3, each of the chord operations shown in Figure 2 is utilized in the four corners of the coarsening region to create the circular chord that is subsequently removed in the third image from the left to produce the coarsened mesh. The far right image shows the final coarsened mesh after smoothing to improve element quality.




Fig. 3. Coarsening a mesh locally can be accomplished by using the operations described in Figure 2 to create circular chords which are subsequently removed. The image on the left shows an initial mesh that is modified using chordal operations to change the topology of the initial mesh to produce a circular chord (left-middle image) which is subsequently removed to produce the coarsened mesh (right-middle image), followed by smoothing (right image).

3.2 Hexahedral Dual Methods

Many of the quadrilateral dual methods have extensions to hexahedral meshes; however, in hexahedral meshes the operations are based on hexahedral sheets and columns. Sheets and columns are topology-based groups of hexahedra that exist in a conforming hexahedral mesh.

Hexahedral Sheets and Columns

A hexahedral element contains three sets of four topologically parallel edges, as shown in Figure 4. Topologically parallel edges provide the basis for hexahedral sheets. The formation of a sheet begins with a single edge. Once an edge has been chosen, all elements which share that edge are identified. For each of these elements, the three edges which are topologically parallel to the original edge are also identified. These new edges are then used to find another layer of elements and topologically parallel edges. This process is repeated until no new adjacent elements can be found. The set of elements which are traversed during this process makes up a hexahedral sheet. Figure 5 shows a hexahedral mesh with one of the sheets in the mesh defined.




Fig. 4. A hexahedral element's three sets of topologically parallel edges.




Fig. 5. A hexahedral sheet: (a) A hexahedral mesh with one sheet defined. (b) A view of the entire sheet.

A hexahedral element also contains three pairs of topologically opposite quadrilateral faces, as shown in Figure 6. Topologically opposite faces provide the basis for hexahedral columns. The formation of a column begins with a single face. Once a face has been chosen, the elements which share that face are identified. For each of these elements, the face which is topologically opposite of the original face is also identified. These new faces are then used to find another layer of elements and topologically opposite faces. This process is repeated until no new adjacent elements can be found. The set of elements which are traversed during this process makes up a hexahedral column. An important relationship between sheets and columns is that a column defines the intersection of two sheets. This relationship is illustrated in Figure 7.




Fig. 6. A hexahedral element's three pairs of topologically opposite faces.




Fig. 7. A hexahedral column: (a) Two intersecting sheets. (b) The column that defines the intersection of the two sheets in (a).

Sheet and Column Operations

Hexahedral sheet and column operations can be used to modify a hexahedral mesh without introducing non-conforming elements. One such operation is known as sheet extraction [23]. Sheet extraction removes a sheet from a mesh

by simply collapsing the edges that define the sheet and merging the two nodes on each edge, as shown in Figure 8. Merging nodes in this manner decreases element density in the vicinity of the extracted sheet and guarantees that the resulting mesh will be conforming.




Fig. 8. Sheet extraction: (a) A sheet is selected for extraction. (b) The edges that define the sheet are collapsed. (c) The two nodes on each edge are merged, which eliminates the sheet and preserves a conforming hexahedral mesh.

Another hexahedral mesh operation that involves sheets is pillowing [28, 30]. Unlike sheet extraction, which removes an existing sheet from a mesh, pillowing inserts a new sheet into a mesh. As demonstrated in Figure 9, pillowing is performed on a contiguous group of simply-connected hexahedral elements which make up a ‘shrink’ set. These elements are detached from the other elements in the mesh, reduced in size, and pulled away from the rest of the mesh, leaving a gap. A new sheet is then inserted into the gap by reconnecting each of the separated node pairs with a new edge. The new sheet increases element density in the vicinity of the shrink set and ensures the preservation of a conforming hexahedral mesh.

A third hexahedral mesh operation is known as column, or face, collapsing [27, 29]. A column is collapsed by merging diagonally opposite nodes in each quadrilateral face that defines the column, as shown in Figure 10. Since a quadrilateral face has two pairs of diagonally opposite nodes, a column can be collapsed in one of two different directions.

As previously mentioned, a column defines the intersection of two sheets. When a column is collapsed, the two intersecting sheets are altered such that they no longer intersect, as illustrated in Figure 11. The paths of the new sheets are determined by the direction of the collapse. Just like sheet extraction and pillowing, the column collapse operation always preserves a conforming hexahedral mesh. In addition, similar to sheet extraction, the column collapse operation decreases element density in the vicinity of the collapsed column.




Fig. 9. Pillowing: (a) A shrink set is defined. (b) The elements in the shrink set are reduced in size and separated from the rest of the mesh. A sheet is inserted to fill in the gap and preserve a conforming hexahedral mesh. (c) The newly inserted sheet.




Fig. 10. Column collapse operation.




Fig. 11. Redirection of intersecting sheets through column collapsing: (a) Two intersecting sheets. (b) The column defining the intersection is collapsed. (c) The two sheets no longer intersect.

4 Quadrilateral Coarsening via Ring Collapse

The previous sections showed that one method of quadrilateral coarsening is to form and extract chords which are confined to the coarsening region. For quadrilateral meshes, the formation of extraction chords is an unnecessary intermediate step. Alternatively, concentric rings of elements, not necessary tied to the dual chords, can be defined and extracted directly without the need of any chord operations. This ring extraction approach provides a more direct and simplified coarsening procedure, and is described in detail in this section.

4.1 Quadrilateral Rings

The first step of a localized quadrilateral coarsening method is to define a region on an existing mesh where element density is to be reduced. We will define this localized region as a ‘coarsening region’. The boundary of this region will be a simply-connected set of edges, E , within the quadrilateral mesh. We further define node set N , as all the nodes defining the edges in E . We now define a ring of quadrilaterals within the coarsening region that share one or more nodes in N . This ring of quadrilaterals will be defined as a ‘coarsening ring’. Fig. 12 shows an example of coarsening rings identified within a coarsening region.




Fig. 12. The image on the left identifies a coarsening region of an unstructured quadrilateral mesh. The image on the right identifies the coarsening rings for the coarsening region.

4.2 Collapsing Coarsening Rings

Given a ring of quadrilaterals, ‘node groups’ are created by identifying edges that are between adjacent quadrilaterals within the ring. The nodes on either




Fig. 13. Final coarsened mesh after several iterations of ring collapses.

end of each of these edges are assigned to a single node group. When a single node is found to be in multiple node groups, the node groups are combined such that a single node belongs in only one node group.

A coarsening ring can be collapsed by combining, or merging, the nodes within each node group into a single node. The location of this combined node is typically the centroid location of all of the nodes in the node group, but may also be moved to the location of constraining geometric entities to maintain the geometric integrity of the mesh through the ring collapse operation.

5 Algorithm for Localized Quadrilateral Coarsening

Using coarsening rings as defined in the previous section, an iterative algorithm for localized quadrilateral coarsening can be outlined as follows:

1. Identify a contiguous coarsening region and a desired level of coarseness.
2. Identify the coarsening rings within the coarsening region.
3. Collapse coarsening rings based on a priority queue until coarseness level is achieved or quality checks prevent further collapses.
4. Perform topology cleanup operations and smoothing to improve final quality.

Each of these steps will be examined in further detail in this section.

5.1 Defining levels of coarseness

Criteria for the amount of coarsening to be performed on a mesh is used as input to the proposed method. Because of topology constraints of a quadrilateral mesh, it may not be possible to precisely achieve a prescribed density. With simplex meshes, iterative local operations can adjust local mesh density very precisely giving a continuous range of mesh density configurations.

Mesh density resulting from quad and hex coarsening, on the other hand, will inevitably be a step-wise function because of the necessity to remove entire chord loops or sheets in a single operation. In spite of this practical limitation, it is however useful to define a target level of coarseness that will drive the amount of coarsening that will be performed. It should be noted, however, that the achieved density will be the average density of the elements in the coarsening region. With the careful selection of concentric coarsening rings and the application of localized smoothing within the region, the density will tend to decrease towards the center of the coarsening region and transition to higher density at the boundary or the coarsening region.

To do this, we first identify a contiguously connected set of quadrilaterals within the mesh referred to as the coarsening region. We then compute the target number of elements to be removed from this region based upon a user defined target element size. The target number of elements can be calculated in one of several ways. We utilize a coarsening factor, F , to determine the final number of elements to be removed from a region, which is calculated based on desired mesh edge lengths within the coarsening region:

$$F = \frac{l_f^2}{l_0^2} \quad (1)$$

where F is the coarsening factor, l_f is the final average edge length specified by the user, and l_0 is the initial average edge length in the coarsening region.

Using the coarsening factor, the number of elements to be removed from a coarsening region can be calculated using the equation:

$$N_{e-r} = E_t - \frac{E_t}{F} \quad (2)$$

where N_{e-r} is the number of elements to be removed, E_t is the number of elements in the coarsening region, and F is the coarsening factor.

Once the number of quadrilaterals to be removed is determined, a 10% tolerance factor (ϵ) is also calculated so that the number of quadrilaterals actually removed is within ($\pm\epsilon$) of the calculated goal. This tolerance factor is limited to being a minimum of three quadrilaterals or a maximum of 50 quadrilaterals. The values of 10%, three and 50 are heuristic values shown to give reasonable results during our experimentation.

5.2 Identifying coarsening rings

The procedure for identification of coarsening rings uses an advancing front method. Starting from the boundary of the coarsening region, a complete loop of quadrilaterals is identified and marked that share mesh edges with the boundary and that are contained within the coarsening region. Successive rings are identified through local adjacency information in the mesh, traversing to unmarked quadrilaterals in the coarsening region to form additional

loops. Each coarsening ring is limited to being a single element in thickness. Because concentric, topologically circular rings, are desired, cases where there are not enough elements to form a closed loop, or where the loop is not topologically circular are identified. These cases are handled by discarding some of the elements from the advancing front. Eventually all of the elements in the coarsening region are advanced on the front completing the ring identification portion of the algorithm. In most cases, the coarsening region is large enough to contain several concentric coarsening rings, however a single loop or disjoint set of loops can also be used. The right panel of Fig. 12 shows the rings developed within the coarsening region. The alternating numbered regions of darker and lighter shaded grey elements show the set of rings. The dark regions not numbered are locations of elements which are not included as rings because they were part of an invalid ring case.

It is recognized that the definition of coarsening rings may be non-unique for any given coarsening region. Multiple, equally valid configurations of coarsening rings may be defined given the same set of quadrilaterals in a single region. In practice, the approach described provides adequate rings for the collapse procedure described here. Further study may be needed to identify an optimal arrangement of rings to achieve better quality quadrilaterals.

5.3 Collapsing coarsening rings

Once the set of coarsening rings has been created, a subset is chosen for removal. Dewey [31] outlines the complete procedure for collapsing coarsening rings, and is briefly presented here for clarity. The projected location of each of the node groups, the location of the nodes after they have been combined, is calculated to enable quality metric calculations for each of the coarsening rings. The rings are then ordered based on projected quality. The projected quality of a coarsening ring is the quality of worst element bounding the coarsening ring, assuming only this coarsening was removed. The projected quality does not take into account the collapse of other coarsening rings, nor does it account for subsequent smoothing. Coarsening rings are selected for removal based on the ordered list unless the ring fails to meet one of two criteria. First, on a given iteration, no two adjacent rings (sharing nodes), are removed. Second, no ring that would cause the total number of elements removed to be greater than the coarsening goal plus the tolerance (ϵ) factor is removed. This strategy has proven to be successful in fully unstructured quadrilateral meshes.

Once the coarsening rings have been selected, the mesh is ready to be coarsened. Each of the selected coarsening rings is collapsed in succession using the following procedure. The nodes in each node group are moved to the projected location and the quadrilaterals that are part of the ring are deleted. As the quadrilaterals are deleted, any edge that is no longer associated with a quadrilateral (i.e. the quadrilaterals on either side of it have been removed) is deleted and the nodes on either end of the deleted edge are merged together.

At corners, a simple collapse (shown in Figure 14) creates a conformal mesh (compare to the operation shown in Figure 3). In the left and center panels of Figure 14, the node groups are the nodes (drawn as circles) connected by darkened edges. The dashed line indicates the coarsening ring of quadrilaterals being collapsed. In the right panel the circled nodes are the locations of the merged nodes in the final mesh.

Table 1. Coarsening rings available for selection.

Ring ID	Ring Quality	Element Count	Selection Order
1	0.0425	90	4
2	0.2184	83	2
3	0.0297	78	6
4	0.1676	71	3
5	0.3107	58	1
6	0.0345	51	5

A simple example is provided here to demonstrate this process. Table 1 gives the identification, quality, and element count of the rings for the mesh shown in Figure 12. The ring quality shown given in the table is the minimum projected Scaled Jacobian of all quadrilaterals adjacent to the ring assuming the ring is removed [31]. The selection order in Table 1 indicates the priority in which rings will be selected for collapse and is based only on the projected ring quality. A ring is deemed acceptable for collapse if, the current number of quadrilaterals to be collapsed plus the number of quadrilaterals in the ring being considered, is less than the goal number of quadrilaterals to be collapsed (plus a specified tolerance factor). Any ring that meets this qualification may be rejected if it is immediately adjacent to a ring that has previously been selected. The goal number of quadrilaterals to be removed for this example is chosen to be 369, with a tolerance of 36. Here, rings 5 and 2 are highest priority for collapse because their collapse will result in a new mesh with the highest projected quality. The other rings in this set are not chosen because they are immediately adjacent to these two rings. The total number of quadrilaterals collapsed in this iteration is 141, significantly less than the goal number of 369 quadrilaterals to be removed. In this situation, as will be explained in Section 5.5, the algorithm will perform additional iterations until the coarsening goal is reached.

5.4 Mesh clean-up and smoothing

Despite efforts to minimize high-valence nodes during the quality assessment of each potential coarsening ring, a few cases remain where high-valence nodes are formed. Furthermore, the collapse of two rings separated by a single layer of elements may reduce element quality due to multiple node projections for




Fig. 14. Collapsing node groups.

quadrilaterals in the non-removed coarsening ring. Additionally, collapsing quadrilaterals along geometric curves may result in elements with low quality which are not improved by smoothing methods. To improve mesh quality, clean-up operations, which change mesh connectivity to create a more structured mesh, as well as smoothing operations to improve the mesh quality by node movement are applied. Most of these final clean-up operations are described in [32, 33].

5.5 Iterations

During a single pass of coarsening by collapsing the rings, it is possible that the desired number of elements to be removed is not reached. Therefore, the steps above are repeated until the coarsening level is reached or quality or topology constraints prevent removal of further elements in the coarsening region. Figure 13 shows the example of this section after it has gone through several coarsening iterations. Note that the region that is completely outside of the coarsening region is identical in both meshes and the elements remain within the coarsening region fit the prescribed goal.

6 Quad Coarsening Examples

In this section, we demonstrate examples of localized quadrilateral coarsening using the coarsening ring collapse operations identified earlier. Changes to mesh quality and the level of coarsening achieved will also be highlighted.

The quality metric used is the scaled Jacobian metric [35] which ranges from -1.0 to 1.0, where a value of 1.0 represents a perfect square while anything below 0.0 is an inverted (non-convex) element (0.0 typically being a triangle-shaped element). For purposes of this study we identify a scaled Jacobian value greater than 0.2 as acceptable for analysis; a scaled Jacobian value between 0 and 0.2 as marginal; and a value less than zero as unacceptable.

6.1 Circular Disk Example

Figure 15 shows the circular cross-section of a cylinder with an associated quadrilateral mesh (the base mesh is shown in the top left corner). In this example, the entire interior of the mesh is chosen as the coarsening region. Figure 15 shows meshes which have been coarsened to factors of 1.5, 4 and 10, respectively.

Table 2 contains the quality compared to the element count for the four meshes. As the level of coarsening increases, the quality of the mesh is typically reduced. The transition between coarse and fine mesh in this example is very rapid resulting in elements with high aspect ratios. The high aspect transitioning is typical in coarsening regions and is the primary reason for quality degradation.

Table 2. Quality comparison on circular disk model before and after coarsening.

	Element Count	Minimum Scaled Jacobian
Initial Mesh	5477	0.67
Factor 1.5	3638	0.66
Factor 4	1347	0.64
Factor 10	572	0.50

6.2 Lever Mesh Example

Figure 16 shows a lever model that has been meshed with shell elements. This figure illustrates localized coarsening on the interior of the mesh. The original coarsening region contains 116 elements and is reduced to 33 following coarsening. Element quality following coarsening remains high and is outlined in Table 3.

Table 3. Quality comparison on coarse lever model before and after coarsening.

	Element Count	Minimum Scaled Jacobian
Elements in coarsening region	116	0.73
Coarsened Mesh	33	0.69

Figure 17 uses the same surface geometry, but with a much finer mesh. The initial mesh is shown on the left and the coarsened mesh is shown on the right. The level of coarsening and the corresponding drop in quality are comparable to those in circle example. Table 4 also contains the quality compared to the element count for the two meshes. In this example the coarsening is localized




Fig. 15. Quadrilateral coarsening on a mesh of a disk at coarsening factors of 0, 1.5, 4, and 10.




Fig. 16. Localized quadrilateral coarsening on a lever mesh. Quad elements selected for coarsening are highlighted

several elements away from the edge of the shell mesh so that the high stress gradients around the edge of the lever are accurately analyzed.




Fig. 17. Quadrilateral coarsening on lever mesh. Quad elements selected for coarsening are interior to the mesh to maintain density near the boundaries

Table 4. Quality comparison on the lever model before and after coarsening.

	Element Count	Minimum Scaled Jacobian
Initial Mesh	11113	0.71
Coarsened Mesh	3261	0.39

7 Hexahedral Coarsening via Localized Sheet Extraction

Utilizing the sheet and column operations described in Section 3.2, the hexahedral coarsening method presented in this section builds upon recent developments in quadrilateral coarsening [29]. While it is true that some quadrilateral coarsening operations can be directly extended to hexahedral coarsening, by themselves, these operations are not always able to prevent changes in element density from propagating beyond the boundaries of a defined hexahedral coarsening region.

7.1 Previously Developed Coarsening Techniques

As illustrated in Section 3.2, sheet extraction decreases mesh density by removing elements from a mesh. Therefore, sheet extraction is a very useful tool for hexahedral coarsening. However, sheet extraction by itself is generally not sufficient when localized coarsening is desired. This is due to the fact that sheet extraction coarsens along the entire path of the extracted sheet. However, sheets are rarely contained entirely within a region that has been selected for coarsening. As shown in Figure 18, extracting a sheet that extends beyond the boundaries of a defined region decreases mesh density in areas where coarsening is not desired. Therefore, before sheet extraction can occur, it is often necessary to modify the mesh in such a way that produces sheets which are contained entirely within the boundaries of a defined coarsening region.




Fig. 18. Global coarsening: (a) A sheet passes through a region selected for coarsening. (b) When the sheet is extracted, mesh density is decreased both inside and outside the defined coarsening region.

As described in Section 3.2, the paths of intersecting sheets can be altered using the column collapse operation. Figure 19, shows how the column collapse operation can be used to create a sheet that is contained entirely within a coarsening region. Such a sheet can then be extracted to coarsen the region without affecting any other part of the mesh.

The coarsening region shown in Figure 19 extends from the top to the bottom of the mesh. Thus, the column collapsed to create a local sheet is contained entirely in the coarsening region, keeping all mesh modifications local. Suppose the coarsening region is modified so that it only extends a few layers from the top of the mesh, as shown in Figure 20. In this case, the column collapse operation can be used twice to produce a sheet that is contained entirely within the coarsening region. However, as seen in the figure, the first collapse operation is performed on a column which extends

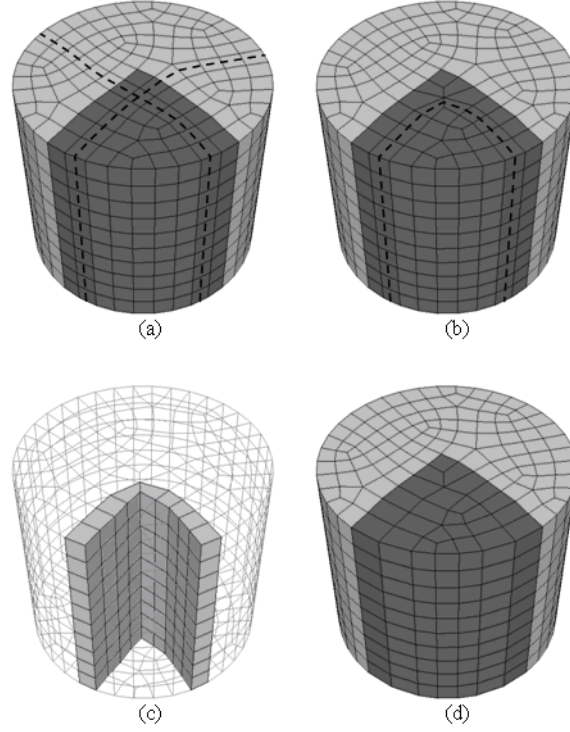


Fig. 19. Localized coarsening: (a) Two intersecting sheets pass through a region selected for coarsening. (b) The column defining the intersection of the two sheets in (a) is collapsed to produce a sheet contained entirely within the coarsening region. (c) The sheet that will be extracted. (d) When the sheet in (c) is extracted, mesh density is only decreased within the defined coarsening region.

beyond the boundaries of the region. Collapsing this column modifies mesh topology and density in areas where coarsening is not desired. This shows that entirely localized coarsening cannot always be accomplished with the column collapse and sheet extraction operations alone. The column collapse process for hexahedral meshes that is explained above, is analogous to the ring collapse procedure for quadrilateral meshes that is described in Section 5.2. Note that in each process, two nodes that share an edge are ultimately collapsed into a single node.

7.2 Entirely Localized Coarsening

The previous examples demonstrate that entirely localized coarsening requires all operations to take place within the boundaries of the selected coarsening region. Referring to Figure 20, it can be seen that the second collapse operation was performed on a column contained within the coarsening region.

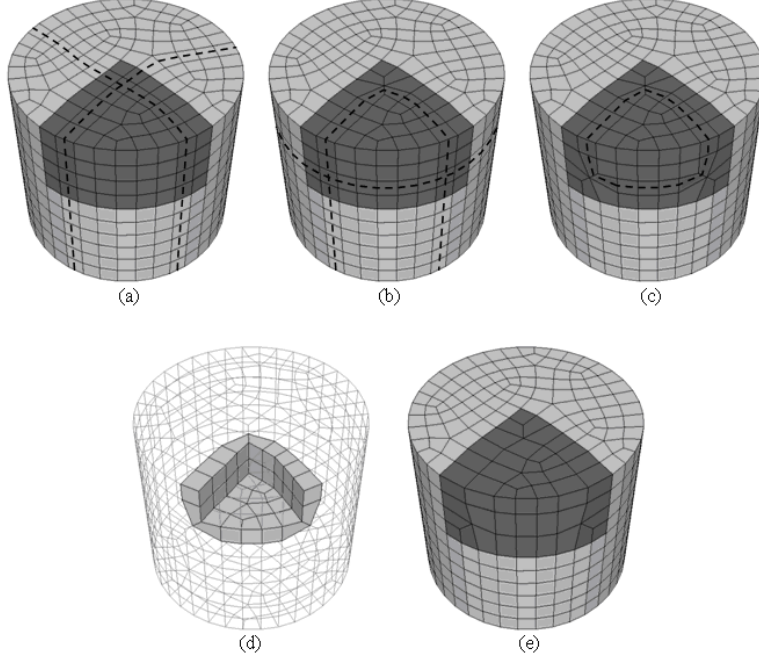


Fig. 20. Semi-localized coarsening: (a) Two intersecting sheets pass through a region selected for coarsening. (b) The column defining the intersection of the two sheets in (a) is collapsed. A sheet formed by the collapse and another intersecting sheet are shown. (c) The column defining the intersection of the two sheets in (b) is collapsed to produce a sheet contained entirely within the coarsening region. (d) The sheet that will be extracted. (e) When the sheet in (d) is extracted, mesh density is only decreased within the defined coarsening region.

Collapsing this column produced a sheet contained within the region without affecting any other part of the mesh. Of course, the formation of this column was accomplished through a previous collapse operation that did affect areas outside the coarsening region. Therefore, a critical aspect of entirely localized coarsening is the creation of local columns. Such columns must be formed in the coarsening region without affecting areas outside the region.

Pillowing of the coarsening region, as the first step in the coarsening process, is a robust and effective way to create local columns without modifying areas outside the coarsening region. As illustrated in Section 3.2, pillowing is a form of refinement because it increases mesh density in the vicinity of the shrink set. For this reason, pillowing is not an obvious solution for coarsening. However, due to the topology constraints that exist in a conforming all-hexahedral mesh, temporarily adding elements is a necessary step when coarsening some regions.

Figure 21 shows how pillowing can be used to produce entirely localized coarsening. Pillowing inserts a new sheet surrounding the coarsening region without modifying any elements outside of the coarsening region. This sheet intersects other sheets that pass through the coarsening region and provides local columns which follow the boundary of the region. Such columns can be collapsed to form sheets contained within the coarsening region without modifying mesh topology or density in areas where coarsening is not desired. These sheets can then be extracted to locally coarsen the region. It should be noted that most of the elements added through pillowing are removed through sheet extraction. Only those elements which are necessary to transition from higher to lower mesh density are left in the mesh. As long as the number of elements removed through sheet extraction is greater than the number of elements added through pillowing, the final mesh density in the coarsening region will be lower than the initial mesh density.




Fig. 21. Entirely localized coarsening: (a) A coarsening region is defined. (b) The sheet that forms when the coarsening region is pillowed. This sheet provides columns which follow the boundary of the region. (c) Collapsing the columns in (b) produces sheets contained entirely within the coarsening region.

7.3 Automated Coarsening Algorithm

For a given region, the process of pillowing, column collapsing, and sheet extraction can be repeated multiple times to achieve various levels of coarsening. The strategy to identify which sheets are to be removed is completely described by Woodbury [34] and briefly outlined here. A simple algorithm has been developed to automate this process for an arbitrary region and level of coarsening. The overall structure of the algorithm is briefly described by the following steps and represented in the flowchart in Figure 22.

1. A coarsening region is defined and a target mesh density for that region is determined based on input given by a user.

2. Every sheet that passes through the coarsening region is found. Sheets contained entirely within the coarsening region are distinguished from those that extend beyond the region.
3. Due to a variety of geometry and mesh topology constraints, each sheet is examined to see if it will facilitate valid collapses and extractions during the coarsening process. Sheets that are unable to facilitate valid collapses and extractions are ignored from this point on.
4. For each acceptable sheet, a shape quality metric [35] is used to estimate how the quality of the mesh will be affected if that sheet, or the portion of that sheet contained in the coarsening region, is extracted. Sheets that will potentially produce a higher mesh quality are given higher priority.
5. If there are any sheets contained entirely within the coarsening region, then valid combinations of those sheets are analyzed. The combination that, when extracted, will produce a mesh density that is closest to the target mesh density without over-coarsening is saved. If no acceptable combination is found, the algorithm moves to step 6. Otherwise, steps 6 through 8 are skipped because no other operations are needed before sheet extraction.
6. If there are any sheets that extend beyond the coarsening region, then valid combinations of those sheets are analyzed. For each combination, two coarsening options are possible, as shown in Figure 23. These two coarsening options are distinguished by which direction the columns are collapsed. The combination that will produce a mesh density that is closest to the target mesh density without over-coarsening is saved. If no acceptable combination is found, steps 7 through 9 are skipped.
7. A sheet is inserted around the boundary of the coarsening region through pillowing.
8. Columns in the pillow sheet are collapsed in directions which were previously determined when the best sheet combination was saved. These collapses form sheets which are contained entirely within the coarsening region.
9. Sheets contained entirely within the coarsening region are extracted.
10. Steps 2 through 9 are repeated until the target mesh density is achieved (within a certain tolerance) or no more valid sheet combinations are found.
11. If coarsening took place, the remaining elements in the region are smoothed to improve mesh quality [36].

8 Hex Coarsening Examples

The following three examples show some results of the automated coarsening algorithm described in Section 7.3. In each example, the goal was to remove 25, 50, and 75 percent of the elements in the region selected for coarsening, while maintaining acceptable element quality. Quality was measured using the scaled Jacobian [35].

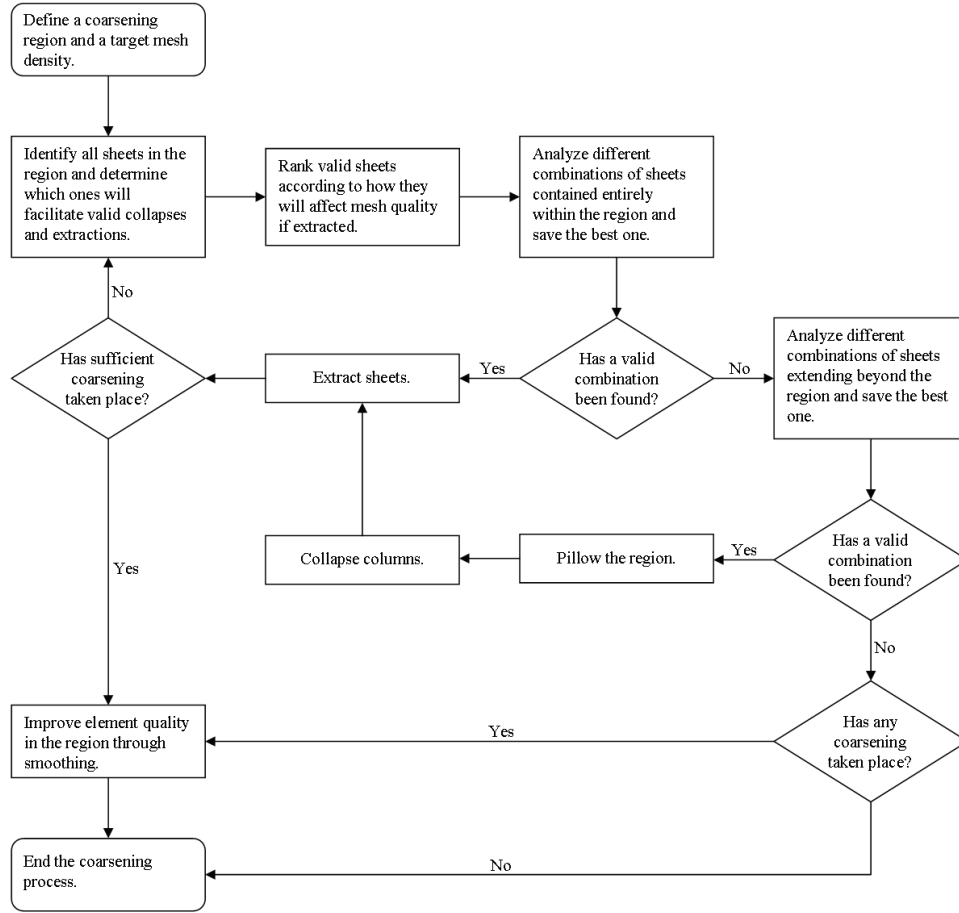


Fig. 22. Flowchart outlining the algorithm for 3D hexahedral coarsening




Fig. 23. Two coarsening options: (a) Columns selected for collapsing. (b) The sheets that will form if the columns are collapsed one way. (c) The sheets that will form if the columns are collapsed the other way.

The first example was performed on a structured mesh of a cube, as shown in Figure 24. The second example was performed on an unstructured multiple-source to single-target swept mesh of a mechanical part, as shown in Figure 25. The final example was performed on an unstructured mesh of a human head generated with an octree based, sheet insertion algorithm [37], as shown in Figures 26 and 27.

Tables 5, 6, and 7 provide element removal, element quality, and coarsening time results for each model. In almost every case, acceptable element quality was maintained and a density that very nearly reflects the target mesh density was achieved.




Fig. 24. Structured cube example: (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

Table 5. Coarsening Results for Cube Model

Target Percent Removal	Elements in Region	Actual Percent Removal	Min. Scaled Jacobian	Coarsening Time (sec)
0	1331	—	1.00	—
25	1056	20.7	0.47	0.7
50	684	48.6	0.41	0.9
75	355	73.3	0.34	1.1


**Fig. 25.** Unstructured mechanical part example: (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

Table 6. Coarsening Results for Mechanical Part Model

Target Percent Removal	Elements in Region	Actual Percent Removal	Min. Scaled Jacobian	Coarsening Time (sec)
0	7641	—	0.77	—
25	5807	24.0	0.59	5.3
50	4057	46.9	0.32	9.6
75	2205	71.1	0.22	12.5


**Fig. 26.** Unstructured human head example (side view): (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

Table 7. Coarsening Results for Human Head Model

Target Percent Removal	Elements in Region	Actual Percent Removal	Min. Scaled Jacobian	Coarsening Time (sec)
0	10080	—	0.48	—
25	7953	21.1	0.29	13.0
50	5129	49.1	0.17	17.9
75	2615	74.1	0.22	22.5




Fig. 27. Unstructured human head example (top view): (a) Original mesh with coarsening region defined. (b) 25 percent coarsening. (c) 50 percent coarsening. (d) 75 percent coarsening.

9 Conclusion

New methods for locally coarsening conforming all-quadrilateral and all-hexahedral meshes have been presented. Using dual-based operations, these methods work on both structured and unstructured meshes and are not based on undoing previous refinement. Automation of the coarsening process has shown promising results. However, more work is needed to ensure acceptable element quality in complex regions and to improve the efficiency of the overall process.

References

1. Biswas Rupak and R. C. Strawn. Tetrahedral and hexahedral mesh adaptation for CFD problems. *Applied Numerical Mathematics*, 26:135–151, 1998.
2. Don Morton and J. M. Tyler. A new 3D adaptive finite element scheme with 1-irregular hexahedral element meshes. In *Proceedings of the 2000 ACM Symposium on Applied Computing*, pages 99–104, 2000.
3. A. Tam, D. Ait-Ali-Yahia, M. P. Robichaud, M. Moore, V. Kozel, and W. G. Habashi. Anisotropic mesh adaptation for 3D flows on structured and unstructured grids. *Computer Methods in Applied Mechanics and Engineering*, 189:1205–1230, 2000.
4. Peter A. Cavallo, N. Sinha, and G. M. Feldman. Parallel unstructured mesh adaptation for transient moving body and aeropropulsive applications. *42nd AIAA Aerospace Sciences Meeting and Exhibit*, pages 6555–6565, 2004.
5. R. Schneiders. Octree-based hexahedral mesh generation. *International Journal of Computational Geometry and Applications*, 10(4):383–398, 2000.
6. Michael Parrish, M. J. Borden, M. L. Staten, and S. E. Benzley. A selective approach to conformal refinement of unstructured hexahedral finite element meshes. In *Proceedings, 16th International Meshing Roundtable*, pages 251–268. Sandia National Laboratories, September 2007.
7. P. Cignoni, C. Montani, and R. Scopigno. A comparison of mesh simplification algorithms. *Computers and Graphics*, 22(1):37–54, 1998.
8. Samuel Silva, Joaquim Madeira, Carlos Ferreira, and Beatriz Sousa Santos. Comparison of methods for the simplification of mesh models using quality indices and an observer study. volume 6492, page 64921L. SPIE, 2007.
9. Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. *Computer Graphics*, 31(Annual Conference Series):209–216, 1997.
10. Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Mesh optimization. *Computer Graphics*, 27(Annual Conference Series):19–26, 1993.
11. K. Bathe. *Finite Element Procedures*. Prentice Hall, Upper Saddle Rive, New Jersey, 1996.
12. Shingo Takeuchi, Hiromasa Suzuki, Fumihiko Kimura, Takashi Kanai, and Kenji Shimada. Subdivision surface fitting using qem-based mesh simplification and reconstruction of approximated b-spline surfaces. In *PG '00: Proceedings of the 8th Pacific Conference on Computer Graphics and Applications*, page 202, Washington, DC, USA, 2000. IEEE Computer Society.

13. A. Cheng and Z. Zhong. Local coarsening for quadrilateral meshes on auto-body tool surface. Journal of Computer-Aided Design and Computer Graphics, 14(1):50–55, January 2002.
14. D. Kwak, J. Cheon, and Y. Im. Remeshing for metal forming simulating - part 1: Two-dimensional quadrilateral remeshing. International Journal of Numerical Methods in Engineering, 53(11):2463–2500, April 2002.
15. C. Choi. Adaptive mesh refinement/recovery strategy for FEA. Structural Engineering Mechanics, 7(3-4):379–391, March/April 2004.
16. G. P. Nikishkov. Finite element algorithm with adaptive quadtree-octree mesh refinement. Australian and New Zealand Industrial and Applied Mathematics Journal, 46(E):C15–C28, 2005.
17. Timothy J. Tautges and Sarah Knoop. Topology modification of hexahedral meshes using atomic dual-based operations. In Proceedings, 12th International Meshing Roundtable, pages 415–423. Sandia National Laboratories, September 2003.
18. Alla Sheffer and Alper Ungor. Efficient adaptive meshing of parametric models. Journal of Computing and Information Science in Engineering, 123:366–375, 2001.
19. C. K. Choi, E. J. Lee, and W. J. Yu. Adaptive mesh refinement/recovery strategy for FEA. Structural Engineering and Mechanics, 17:379–391, 2004.
20. Yannis Kallinderis and C. Kavouklis. A dynamic adaptation scheme for general 3D hybrid meshes. Computer Methods in Applied Mechanics and Engineering, 194:5019–5050, 2005.
21. Benjamin S. Kirk, J. W. Peterson, R. H. Stogner, and G. F. Carey. libmesh: A c++ library for parallel adaptive mesh refinement/coarsening simulations. Engineering With Computers, 22:237–254, 2006.
22. R. Taghavi. Automatic, parallel and fault tolerant mesh generation from CAD. Engineering With Computers, 12:178–185, 1996.
23. Michael J. Borden, Steven E. Benzley, and Jason F. Shepherd. Coarsening and sheet extraction for all-hexahedral meshes. In Proceedings, 11th International Meshing Roundtable, pages 147–152. Sandia National Laboratories, September 2002.
24. Steven E. Benzley, N. J. Harris, M. A. Scott, M. J. Borden, and S. J. Owen. Conformal refinement and coarsening of unstructured hexahedral meshes. Journal of Computing and Information Science in Engineering, 5:330–337, 2005.
25. Peter J. Murdoch and Steven E. Benzley. The spatial twist continuum. In Proceedings, 4th International Meshing Roundtable, pages 243–251. Sandia National Laboratories, October 1995.
26. Peter J. Murdoch. The Spatial Twist Continuum: A Dual Representation of the All Hexahedral Finite Element Mesh. Published Doctoral Dissertation, Brigham Young University, December 1995.
27. Jason F. Shepherd. Topologic and Geometric Constraint-Based Hexahedral Mesh Generation. Published Doctoral Dissertation, University of Utah, May 2007.
28. Karl Merkley, Corey D. Ernst, Jason F. Shepherd, and M. J. Borden. Methods and applications of generalized sheet insertion for hexahedral meshing. In Proceedings, 16th International Meshing Roundtable, pages 233–250. Sandia National Laboratories, September 2007.

29. Matthew L. Staten, S. E. Benzley, and M. A. Scott. A methodology for quadrilateral finite element mesh coarsening. Engineering With Computers, 24(3):241–251, 2008.
30. Scott A. Mitchell and Timothy J. Tautges. Pillowing doublets: Refining a mesh to ensure that faces share at most one edge. In Proceedings, 4th International Meshing Roundtable, pages 231–240. Sandia National Laboratories, October 1995.
31. Mark W. Dewey. Automated Quadrilateral Coarsening by Ring Collapse. M.S. Thesis, Brigham Young University, 2008.
32. Scott A. Canann, S. N. Muthukrishnan, and R. K. Phillips. Topological improvement procedures for quadrilateral finite element meshes. Engineering with Computers, 14:168–177, 1998.
33. Paul Kinney. Cleanup: Improving quadrilateral finite element meshes. In Proceedings of the 6th International Meshing Roundtable, pages 449–461, October 1997.
34. Adam Woodbury. Localized Coarsening of Conforming All-Hexahedral Meshes. M.S. Thesis, Brigham Young University, 2008.
35. Patrick M. Knupp. Algebraic mesh quality metrics. SIAM J. Sci. Comput., 23(1):193–218, 2001.
36. Patrick M. Knupp. Hexahedral and tetrahedral mesh shape optimization. International Journal for Numerical Methods in Engineering, 58(1):319–332, 2003.
37. Jason F. Shepherd, Yongjie Zhang, Claurissa Tuttle, and Claudio T. Silva. Quality improvement and boolean-like cutting operations in hexahedral meshes. In Proceedings, 10th Conference of the International Society of Grid Generation, September 2007.